Introducing the Ethernet *Sniffer*

The Sniffer™

Network General

DISCLAIMER OF WARRANTIES

The information in this document has been reviewed and is believed to be reliable; nevertheless, Network General Corporation makes no warranties, either expressed or implied, with respect to this manual or with respect to the software and hardware described in this manual, its quality, performance, merchantability, or fitness for any particular purpose. The entire risk as to its quality and performance is with the buyer. The software herein is transferred "AS IS."

Network General Corporation reserves the right to make changes to any products described herein to improve their function or design.

In no event will Network General Corporation be liable for direct, indirect, incidental or consequential damages at law or in equity resulting from any defect in the software, even if Network General Corporation has been advised of the possibility of such damages. Some states do not allow the exclusion or limitation of implied warranties or liability for incidental or consequential damages, so the above limitation or exclusion may not apply to you.

Manual prepared by Paul Berry
Appendices by Leonard J. Shustek
April 1987

Introducing the Ethernet *Sniffer*

The Sniffer collects, analyzes and interprets data transmitted in a local-area network. It exists in two versions: one for Ethernet, the other for the Token Ring. The versions provide the same capabilities, and report them in the same way. However, they aren't identical because differences between the two systems of transmission require corresponding differences in what the Sniffer does.

You have the Ethernet Sniffer. You have a complete machine, with the software already installed. However, you don't yet have a complete Ethernet version of the Sniffer's *Operation and Reference Manual.* That's because (as usual) the engineers are ahead of the documentation. So, to tide you over, the document you're now reading gives you:

1. An overview of the Ethernet Sniffer. This material is much like Chapter 1 of the Token Ring Sniffer's *Operation and Reference Manual.*

2. A list of ways that the Ethernet Sniffer differs from the Token Ring Sniffer, so you can make reasonably effective use of the Token Ring version of the manual. In particular, this section tells you where the Token Ring manual applies to you too and where it doesn't.

Overview

The Sniffer attaches to a network in the same way as other stations. There it listens to the traffic that is broadcast on the Ethernet bus. It maintains real-time counts of the flow of frames. It selects some of the frames its sees, and records them for later analysis. Its detailed reports include protocol interpreters that translate the various levels of code and display them in English, from the data link level on which everything else rests, up to the session level used by network applications.

With its detailed records of exactly what transpires during network transactions, it is a powerful tool for trouble-shooting and tuning a network, and for testing and refining high-performance network software.

The Sniffer Is Self-Contained

The Sniffer is a fully portable computer and is completely self contained. It comes with its own Ethernet adapter already installed, its own hard disk, and its own operating system and software ready to run.

To start using the Sniffer, you need only plug its power cord to a suitable outlet and connect it to the network. You may be able to plug it directly to an existing tap (perhaps temporarily replacing the station usually connected there). Alternatively, you may wish to install a new tap just for the Sniffer. Because there is such a variety of possible connectors, the Sniffer does not include the cable by which you attach it to the network, but this document explains what you'll need and how you attach it.

The only customizing of the software you'll probably find desirable is to augment the Sniffer's definition file of station names. The Sniffer can then display both the hardware addresses it observes and the names by which you refer to the various machines. (You can build your name tables as you go along; see *Managing Names*.)

Menu-Driven Controls

An autoexec batch file already installed on the hard disk starts the Sniffer software as soon as you turn the machine on. You operate the Sniffer from a menu screen. You move the cursor to the choices you want, select options by pressing the space bar while they're highlighted, and press *Enter* or one of the function keys to start an action. Whenever a function key is operative, it's highlighted and labeled in the screen display.

There is no command language, and there are no commands to learn. About the only information you'll supply by typing is the name for a file you wish to save.

When you exit from the Sniffer's software, it returns you to its operating system, COMPAQ DOS 3.10. The Sniffer is then a

standard AT-class personal computer operating under DOS. A DOS manual is included with the Sniffer.

Color Monitor or LCD Display

The Sniffer's built-in monitor is high-resolution 8-level grey-scale monochrome, with a green phosphor. You can also connect your own color monitor, or an external LCD display. The Sniffer provides a DB-9 jack for an RGBI monitor that supports the IBM color graphics monitor, and an RCA jack for composite video. You have only to plug in your equipment and report that it's there when the Sniffer program asks.

The Sniffer Is a Specialized Station on the Network

Like most of the other network devices, the Sniffer is an independent computer with its own software and hardware, and its own network adapter. It does not need (and does not include) a copy of the network management software used by ordinary stations on the network.

As far as the other stations on the network are concerned, the Sniffer is a passive member. Like any ordinary station on the network, it hears the transmissions from all other stations. It notes them for analysis, but it never responds to the other stations. It never acts as receiver for messages that other stations send. It originates traffic addressed to other stations only in a test mode designed to load the network. It can emit test pulses to look for cable faults.

The Sniffer Hears Every Frame

Ethernet is a bus system. All stations are connected to a common bus. Like every other station on the network, the Sniffer hears every frame transmitted.

The Sniffer's adapter card makes a temporary record of each frame, and passes it to the Sniffer's on-board processor for review. The processor filters these just-received frame. It records those that pass the *capture filter* you've set, and records them for later analysis and interpretation. It discards the rest.

Capture Filters

The number of frames reaching the Sniffer's adapter is potentially so large that it's essential to select only a subset. The Sniffer applies a filter to each newly-arrived frame, and discards the frames that do not meet its test. Capture Filters are of three types:

- ***Selection by station address:*** include frames sent from or received by a particular station or pair of stations.

- ***Selection by protocol:*** include frames containing any of the protocols you specify (using either Ethernet or IEEE 802.3 protocol descriptions).

- ***Selection by pattern:*** include frames containing a specified pattern of data at a particular position in the frame.

(For example, a typical filter might admit only messages to or from a particular user and a server with which the user is experiencing a problem, and only those frames involving a particular protocol).

Setting an appropriate filter is your first step in collecting data. Often, the majority of arriving frames are immediately discarded. The frames that your filter admits then pass to a buffer area, from which you may display them, send them to storage, or discard them.

Real-Time Meters and Counters

While the Sniffer is collecting data, it measures the rate at which frames are arriving, and gives you a real-time graphic display of *meters* (which show the data-rate), and *counters* (which show a running total of the numbers of frames transmitted).

You can display the traffic density as kilobytes per second or as frames per second. For either, you can elect to show them as absolute values or as percentages of the network's available bandwidth, and on a linear or on a logarithmic scale.

Counters can tabulate frames by destination or by source, or cross-tabulate them by station pairs. The display is expanded in real time. As the Sniffer notices traffic involving stations it hasn't seen before, it makes room in the display to include them.

The Capture Buffer

After they've been counted, frames that the filter accepts pass to the Capture Buffer. (On the way, they're examined by the Trigger Detector, described in a moment.) The Capture Buffer has room for a moderate number of frames (hundreds of medium-sized frames, or thousands of minimal-sized ones). Frames accumulate in the Buffer in the order they are received.

When the Capture Buffer becomes full, the Sniffer may halt capture, or may discard older frames to make way for new arrivals, as you've elected. If you do nothing to retain the frames in the Capture Buffer, the Sniffer automatically discards them; in that case, the frames that remain in the Buffer are the ones most recently received.

The Trigger Detector Scans Incoming Frames

The Trigger Detector scans the stream of incoming frames. It's located after the Capture Filter, so that it looks only at frames that have passed through the filter but haven't yet reached the Capture Buffer (see Figure 1-1).

The Trigger Detector looks for a frame containing a particular pattern that you've described. When it finds such a frame, it signals a *trigger event.* The trigger event freezes the Capture Buffer so you can examine the trigger frame and the frames that precede or follow it.

A Trigger Event Stops Collection and Freezes the Buffer

When the Trigger Detector signals a trigger event, capture ceases, either immediately or with enough delay to collect some of the following frames. Once capture has been halted, you can:

- Copy the contents of the Capture Buffer to a file for later analysis or display.
- Browse through various displays of the frames in the Capture Buffer.
- Impose a display filter to select which frames are displayed.
- Select one or more *views* (ways of displaying a frame).
- Print the contents of the buffer, according to the filters and views you've specified.

A trigger event halts the processing of incoming data. It causes the Sniffer to cease capturing frames until you say you're again ready to receive them.

Specifying the Trigger Pattern

A trigger pattern is a set of characters at a particular position in a frame. You can make the test match either the presence or the absence of the pattern.

For example, if you're examining complaints of intermittent problems with access to a particular file server, you might set up a collection filter that accepts only frames to or from that station, and a trigger that signals when it spots an error return code.

The frame that matches the trigger pattern is called the *trigger frame.* When it appears in your display of the Capture Buffer, the trigger frame is identified by a letter T beside it. One of the actions during display is "Jump to trigger frame."

Frames Surrounding the Trigger Frame

When you set up a trigger pattern, you also indicate where in the Capture Buffer you want the trigger frame to appear. That determines whether the Buffer contains frames that preceded the trigger frame, frames that followed it, or some on either side.

Displaying the Frames in the Capture Buffer

You have many options for displaying the contents of the Capture Buffer, either at the Sniffer's screen or to a printer. (You can direct printer output either to a locally-attached printer, or to a file on one of the Sniffer's disk drives.)

You can set up a display filter so that frames which don't interest you are omitted from the display (even though they remain in the Capture Buffer). The mechanism for filtering frames from the Capture Buffer is like the mechanism for filtering frames during capture.

Saving the Capture Buffer for Later Analysis

From the keyboard, you can select a command that saves the contents of the Capture Buffer to a file. You can save the entire Capture Buffer, or just the frames that are selected by your current Display Filter.

All displays and analyses work with the data in the Capture Buffer. You can display data that has arrived and is still in the Buffer, or you can load the Capture Buffer with data you earlier saved to a file.

Selecting the Form of Display

The display may contain any or all of the following three reports:

- ***Hexadecimal view.*** The entire frame is listed. You can elect whether you want character data displayed according to ASCII or EBCDIC conventions.

- ***Detail view.*** Each frame is decoded to show the type of frame and the values of its various fields. If you provide a file of symbolic names for station addresses, the *detail* view augments the station names with the symbolic names provided in your file of definitions.

  For high-level frames, the interpretation may take several levels. The "higher" level interpretation of a frame is more deeply nested within it. The various interpretations are shown with the "higher" protocol levels (i.e. the ones that are more "deeply" embedded) after the lower ones.

- ***Summary view.*** This condensed form abbreviates and truncates some of the information from the hexadecimal

view and some of the information from the *detail* view. It forces each level of interpretation to fit on a single line. The display contains one line for each level of protocol in the frame. You can elect to show only the highest level; in that case, the *summary* view has one line per frame.

The *hexadecimal* view and the *detail* view show data for just one frame. The *summary* view shows not only the frame you're now examining, but a few on either side of it as well, to give context.

Windows in the Display

Each view you elect appears in a window. The screen is divided into one, two or three equal-sized windows, one above the other, according to the number of views you request.

The window that contains the cursor bar is the active window. The *Tab* key moves the highlight from one window to the next, activating the window where it arrives. When a frame's display won't fit within its window, you can scroll the active window to see the information you want. You can also zoom in to the active window, temporarily, giving it the entire screen until you zoom out and restore the other windows.

Two-Station Display

Frequently, analysis concerns the flow of commands back and forth between a pair of stations. In that situation, it is often helpful to elect *two-station display*. Frames from one station are shown on one side of the screen or paper, frames from the other on the other side. (Frames that are not part of the two-way interaction are also shown, but in the default format.)

Dual Viewports

Sometimes it's important to compare a frame from one part of the Capture Buffer with a frame that arrived earlier or later. You can do that by electing *dual viewports*. The screen is split into left and right halves. In each window, you can scroll the two sides independently, permitting you to concentrate on one frame on the left and another frame on the right.

Saving and Restoring Setups

Because there's a rich choice of options concerning what to display and how to display it, the Sniffer lets you save a record of the way you are filtering and displaying frames, so that you can readily restore the setup at a subsequent work session.

The Protocol Interpreters

The Sniffer doesn't just capture and store frames from the network. It also interprets them. When you select the *detail* view, for each frame you get a set of interpretations, one interpretation for each level of protocol that the frame contains. The interpreter labels and decodes the standard fields in each frame, making it easy to see the message conveyed.

When you select both a *detail* view and a *summary* view, the Sniffer automatically scrolls the *detail* view so that the interpretation shown there matches the level you've highlighted in the *summary* view.

If your network transmits frames whose protocol is unknown to the Sniffer's interpreters, it's possible to augment Network General's interpreters with a custom interpreter of your own. To write one, you'll need detailed familiarity with the protocol, with the network's DLC and LLC conventions (data link control and logical link control), and the C programming language. Specifications for such an interpreter are provided in Appendix C to the *Operation and Reference Manual.*

Traffic Generator

The Traffic Generator permits you to load the network with background traffic messages sent from your Sniffer. You can specify the addressee, the total length of the frame and the interval between frames in milliseconds. You can also specify the content of the first thirty bytes, to control how other stations will treat the frames the Sniffer generates.

You can generate frames of any length between 60 and 1514 bytes (the least and greatest lengths for a legal Ethernet frame), or of lengths between 12 and 59 bytes (which other stations will treat as collision fragments).

While the Sniffer is generating traffic, that occupies it completely so it cannot at the same time perform its other functions.

Playback

From any of the saved capture files, you can regenerate the display of meters and counters, just as though the packets were arriving from the network in real time. Playback permits you to experiment with the display options that are available during capture even when you're not capturing "live" frames. It also makes a convenient way to demonstrate a particular sequence to your colleagues, or to introduce the Sniffer to people who haven't used it. At a tutorial session based on playbacks, you can:

- Observe the dynamics of capture, including changes in traffic density and the accumulation of traffic counts.

- Experiment with capture filters to weed out extraneous frames.

- Experiment with triggering (that is, freezing the capture buffer when a particular trigger pattern is received).

- Experiment with recording only part of each captured frame.

To run a playback, you don't have to be connected to the network, so you can use it for experiment, training or demonstration even when the network is down, or you're not connected to it.

Schematic View

Figure 1-1 conceptualizes the Sniffer's various functions. It's more a cartoon than a formal diagram, but it correctly conveys the flow of data in the Sniffer.

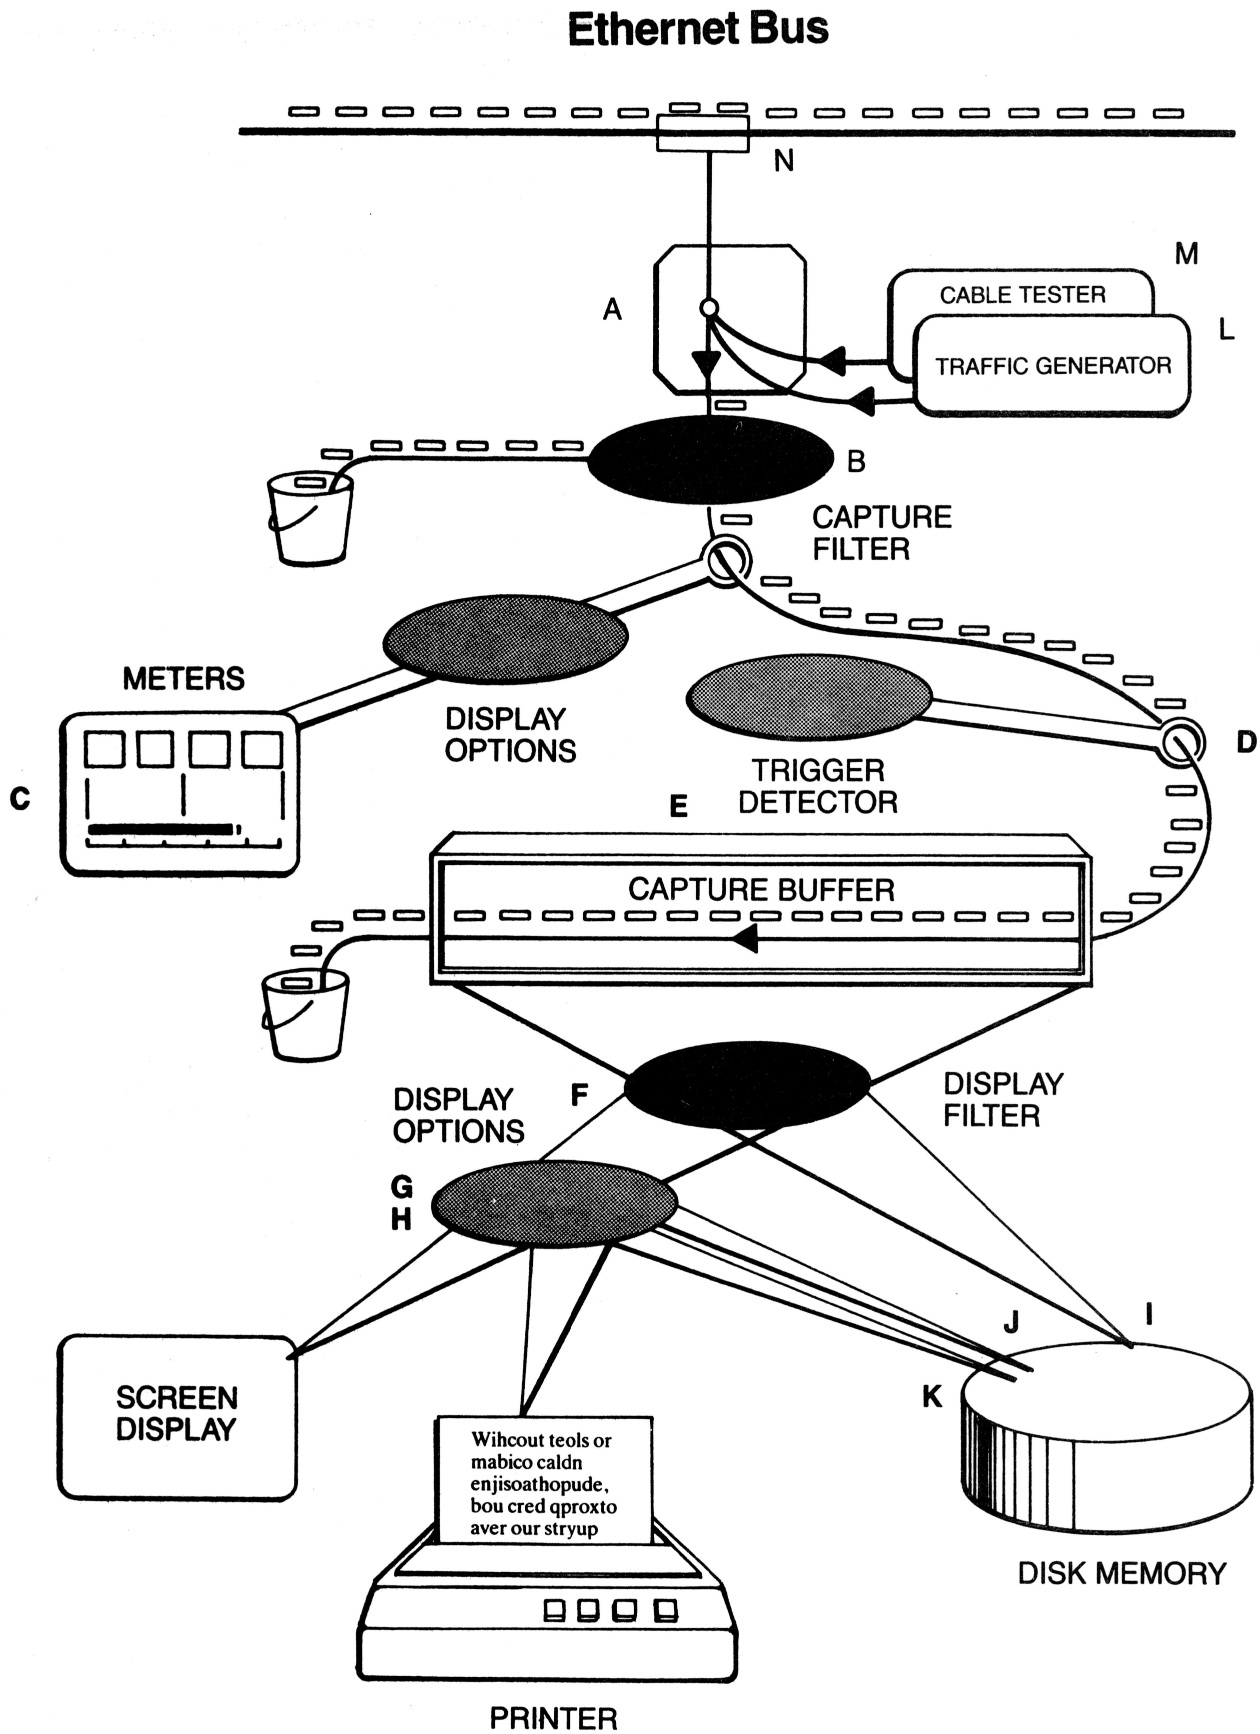

Figure 1-1: Schematic representation of the Sniffer's functions

Key to Figure

A. At the top, the Sniffer is connected to the network's bus, on which all transmissions are broadcast. The Sniffer's specially-modified network adapter card retains a copy for the Sniffer's analysis.

B. The Capture Filter immediately discards frames that don't meet its address and protocol filters.

C. During capture, the meters and counters provide real-time display of activity of the captured frames.

D. The trigger scans the accepted frames for patterns for which it has been alerted.

E. The capture buffer holds frames captured from the network, or loaded from a file. Unless capture is halted manually or by a trigger, as the buffer is filled by arriving frames, the earlier frames are discarded to make room for the new ones.

F. Display filters select the frames visible on the screen or printer display.

G. Display options select the number and type of views displayed.

H. The frame interpreters decode fields in the frame being viewed.

I. The set of frames in the capture buffer can be saved to a file (with or without winnowing by the display filter).

J. Set-ups (display filters, view options, etc.) may also be saved to a file.

K. A stored file of definitions helps interpret station names in the display.

L. The traffic generator transmits frames containing the data you've specified at the interval you've specified.

M. The cable tester emits a pulse and times its echo, looking for the reflections characteristic of a short or an open cable.

N. Vampire tap to the Ethernet bus.

What to Add and What to Ignore from the Token Ring Version of the Manual

Because the two versions of the Sniffer share not only the same philosophy but also a fair amount of code, much of the Token Ring version of the manual is equally applicable to the Ethernet Sniffer. What follows points out where they differ, and what you'll need to know to use the Sniffer on Ethernet.

In what follows, the notation *TR-* followed by a number indicates a chapter from the Token Ring manual. For example, TR-2 indicates "Chapter 2 of the Token Ring manual;" *TR-Add* indicates the *Addendum* to the Token Ring manual.

TR-1: Overview

Ignore it. The Overview you've just read essentially duplicates and replaces it.

TR-2: The Sniffer at Work

It describes some case histories taken from the Token Ring. Ignore the details. You may want to skim them for general impressions on the way the Sniffer can be used.

TR-3: Setting Up the Sniffer

The general description of unpacking the Sniffer is correct. On page 41 TR-3 mentions the card protector in the floppy drive. It should also say that you should hold on to the card, because you'll need it if you have to travel with the Sniffer or ship it back to Network General for repair.

Ethernet Connection

Naturally, the Ethernet Sniffer does not have the Token Ring connector described on page 42. However, it does have an adapter plate, which you may need to secure your Ethernet cable to the Sniffer's adapter card.

The adapter in the Ethernet Sniffer is different from the one shown in Figure 3-1, page 43 but it is still mounted in slot 3 (in the slot otherwise used for the Token Ring adapter). It has two connectors. At the top is a 15-pin female D-connector. Below it is coaxial connector covered with a red cap. *Use only the D-connector.* There is at present *no* use for the coaxial connector. Don't remove its cap.

Cable

To connect the Sniffer to the network, you will need a cable with a 15-pin male D-connector to mate with the connector on the Sniffer's adapter.

At the other end, the cable must connect to a transceiver on your Ethernet bus cable. You'll need whatever cable or connector your transceiver requires. You may already have such cables for other stations on your network, or you may have to order one. A common but not universal arrangement is to have a male 15-pin connector on the transceiver. In that case your cable will need 15-pin D-connectors at each end, one male, one female. These cables do not come in standard lengths; nor is the length critical, provided you don't exceed 50 meters. Obtain a cable that reaches conveniently from the transceiver to the Sniffer. (You'll need to know the length of this and other network cables in order to interpret information reported by the cable tester.)

Lockposts vs. Screws

Ethernet cables are commonly secured by a slide on the cable connector which attaches to a lockpost on the device. Personal computers, including the Sniffer, are generally provided with screw posts that secure cables by screwing them down. If your cable is designed for lockposts, in order to secure it to the Sniffer's adapter card, you will need to install an adapter plate on the end of the cable that attaches to the Sniffer. The adapter plate clips on to the plug, and provides screws by which the plug can be secured. The adapter plate is included with the Sniffer. If you don't need it now, set it aside for some future occasion, and skip the paragraph that follows.

Installing the Adapter Plate for Screw Connections

To do this, you'll need a small flat-bladed screwdriver.

Slide the threaded clips onto both ends of the adapter plate, and insert the screws into the clips. In Figure E-2, at the top you can see one of the clips positioned ready to slide on, and at the bottom a clip in place, with the screw inserted.

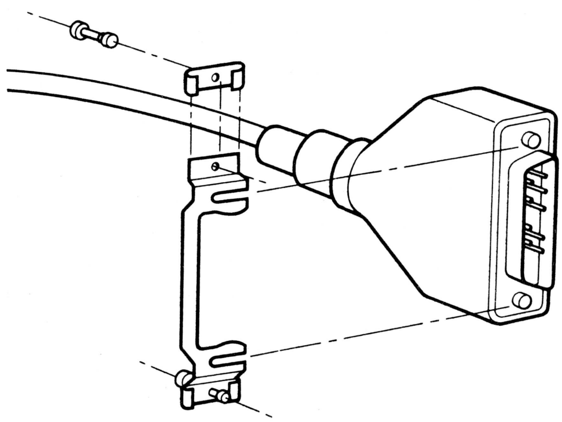

Figure E-2: Adapter plate ready for attachment to a D-connector with lockposts.

Align the slots in the adapter plate with the indents in the lockposts on the transceiver cable. Press the adapter plate until it snaps into position on the connector.

Plug the connector with its adapter plate to the DB-15 connector in the Sniffer's slot 3. Fasten the screws to the threaded receptacles above and below the connector, as shown in Figure E-3.

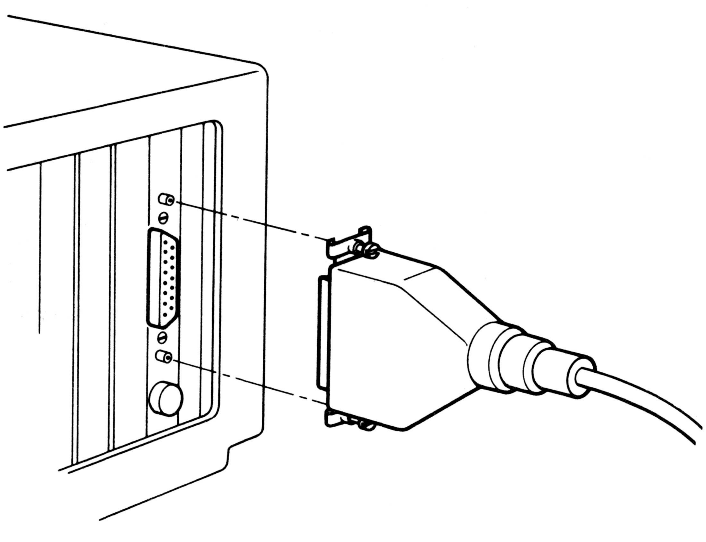

Figure E-3: Connecting a cable with adapter plate to the Sniffer's network adapter card.

Installing a Transceiver

If there is no available transceiver at which you can connect the Sniffer to the network, you may have to obtain and install a new one. Network General uses the NT100 transceiver manufactured by MICOM-Interlan; others may be equally suitable.

The transceiver is a passive repeater, which acts as the point of connection between one station and the network bus. Basically, the transceiver attaches a "T" to the network bus. While other transmissions pass straight through, the branch leads through the transceiver's electronics to the cable attached to a station on the network.

Some models of transceiver require you to cut the network bus, and connect the severed ends to either side of the transceiver by a coaxial N-connector, or a BNC connector if the network uses thin Ethernet cable. Others, called "vampire taps," leave the bus intact and make their connection by piercing the cable's insulation. The vampire connection may be easier to install, with less risk of additional reflection from the connectors.

LCD Option

In addition to the color monitor described on page 43 of TR-3, you may also connect an LCD display. See TR-Add page 171.

Backup

The instructions about backing up the software on page 40 of TR-3 are generally correct. However, the utility MAKEFLOP is no longer provided. To make backup copies, see the description of the BACKUP command in the COMPAQ DOS manual.

Main Menu

The Sniffer's Main Menu is much as shown on page 50 and described on page 52 of TR-3, except that it also has the options *Cable Tester* (not shared with the Token Ring) and *Traffic Generator* (added since figure 3-2 was printed).

Of course, in its *Capture* menu, the Ethernet Sniffer shows the option *From <Ethernet>* rather than *From <Token Ring>*.

TR-4: Capturing Frames

You can capture frames either live from Ethernet, or by replaying them from a file of frames captured earlier. Page 55 of TR-4 doesn't mention the playback option (because it has only just been added); it is described in TR-Add, at page 156.

Insertion

Ethernet has nothing comparable to the procedure required to insert a station into a Token Ring. Ignore the second half of page 55.

Network Types

There exist two systems for identifying the protocol used in a frame transmitted over Ethernet. The first is the one originally introduced by Xerox. It uses two bytes to designate the frame's type, referred to as its *Ethertype*. A second classification was introduced with IEEE 802.3. It starts with a two-byte length, followed by an 802.2 LLC header. Usually (but not necessarily) all stations on a network use one or the other but not both.

To indicate how the Sniffer should interpret frames during display, the *Display/Summary* menu includes a two-position radio control (that is, a choice between mutually exclusive alternatives). You can set the control to indicate *Ethertype* format, or the length-and-header format of IEEE 802.3. Setting this switch affects the Sniffer at two points:

- ***Before capture:*** The switch's setting determines not only the interpretation of captured frames, but also which options you'll see in the menu of protocols for setting the capture filter. (The default setting is *Ethertype*, which is at present more common.)

- ***During display:*** The switch also determines how the Sniffer interprets the captured frames that it displays.

Whenever frames are captured, or loaded into the Capture Buffer from a saved file, the Sniffer sets the switch itself, based on its inspection of the captured frames. Because most Ethertype codes have values higher than the length of the longest frame, it is usually possible to deduce which protocol the captured frames exhibit. The Sniffer sets the switch to *Ethertype* when it finds that for most frames the value of the first two bytes is greater than the actual length of the frame or equal to hex `0200` (the PUP Ethertype).

You may override the Sniffer's judgment about the type of frames it has received by setting the switch yourself after you've completed capture.

If you use an inappropriate setting *before* capture, you may set up a capture filter that will look for protocols that aren't present, and thus select few if any frames. If you use an inappropriate setting *during display*, the Sniffer may (wrongly) report many frames that have "unknown Ethertype" or "unknown SAP" and "invalid length." That's easily remedied by setting the switch back as it should be and returning to the display.

Capture Filters

The description of capture filters is generally correct, apart from the fact that TR-4 doesn't mention the choice between *Ethertype* and IEEE 802.3. Naturally, the Ethernet Sniffer shows you a menu appropriate to Ethernet, rather than the Token Ring options that are illustrated in Figure 4-4.

Protocols in the Capture Filter

As shown in Figure 4-4 of TR-4 (page 59) you may select from a menu of low-level protocols for capture. The actual menu, of course, is one appropriate to Ethernet, rather than the list of Token Ring protocols shown in the figure.

Defective Frames Filter

During capture on Ethernet, you may elect to filter to include or exclude various types of defective frames, by checking of those you will accept from the following list:

- Good frames.
- Frames with bad CRC (cyclic redundancy check).
- Collision fragments (frames shorter than 60 bytes).
- Frames with bad alignment (that is, received with a length not a multiple of 16 bits).

These options do not exist in the Token Ring environment, and are not mentioned in TR-4. They appear below *Pattern match* in the menu shown on page 60 of TR-4.

Pattern Match

The procedure for setting a pattern as part of the capture filter is essentially the same as described in TR-4, pp. 60-61. However, Ethernet does not require the distinction between "frame relative" and "data relative" locations for the pattern, and so those choices (visible in Figure 4-5 of TR-4) don't appear in the Ethernet version.

Trigger

The procedures are identical to those in the Token Ring version (TR-4, pp.62-67) except that there is no need to indicate whether a trigger pattern is frame-relative or data-relative, so that choice (visible in Figure 4-10) does not appear in the Ethernet menu.

Capture

The description in TR-4 pp. 67-73 is correct as far as it goes. It does not mention capture by playing back a file of previously captured frames. For that, see TR-Add, page 156-158.

When you elect to capture from a file of captured frames, the files from which you select must have names which include the extension .ENC (for "Ethernet capture") rather than the .TRC extension shown on page 157 of TR-Add. At the bottom of page 158, the relevant choice is of course "<Ethernet>" rather than "<Token Ring>."

The Sniffer now permits you to capture only the first so-many bytes of a frame. The procedure is described on TR-Add page 159, and its effect on page 160.

TR-5: Displaying and Interpreting Captured Data

The display facilities described in Chapter 5 are all equally applicable to the Ethernet version. The display facilities start with display filters, and offer three views of each frame, as described on page 82.

There are a few obvious differences. The Ethernet display filters permit you to force Ethertype or IEEE 802.3 interpretation of frame types. (If you don't make a choice, the Sniffer deduces which is likely. The menu option permits you to override that.)

The list of protocols you may select (shown in Figure 5-5, page 81) is one appropriate to Ethernet and differs from the list shown in Figure 5-5.

ASCII vs EBCDIC

While the Token Ring Sniffer is willing to guess whether characters in the Hexadecimal view should be transliterated as ASCII or EBCDIC (*Dynamic mode* in Figure 5-7), the Ethernet version requires you to choose one way or the other.

Address Recognized and Frame Copied

These are described in TR-5 on page 85. They have no counterparts in the Ethernet version. Ignore them.

Names in the Summary View

Since TR-5 was printed, the treatment of station names has been expanded, so that each name is now treated as a pair, consisting of the name of he station and the protocol level in which the name can occur. This change is described in TR-Add 161-164, and applies equally to the Token Ring and the Ethernet versions of the Sniffer. This change affects the details of *Managing names* (TR-5 pp. 99-102).

Resolve Names

Since publication of the Token Ring manual, the ability to look up station address encountered on the network in an external dictionary of names has been added both to the Token Ring and to the Ethernet Sniffer. The procedure is described on pages 165-166 of TR-Add. For the Ethernet version, the reference file must have an extension of `.END` rather than the `.TRD` mentioned on page 165.

Address-Level Filters

The qualification of names based on the protocol on which they occur makes possible *Address Level Filters.* These display filters pass or reject a frame based on whether it includes an address at one of the protocol levels you select. (The selection is not based on what the address is, but on the protocol level at which it occurs.) This change applies equally to the Ethernet and to the Token Ring versions; it is described in TR-Add 167-168.

Printing or Saving Only Part of the Capture Buffer

The Buffer now permits you to save to a file to print a display based on a range of positions within the Capture Buffer. This amends the description of Printing (TR-5 pp. 102-3) and Saving (TR-5, p. 104) by the material found at TR-Add 172-173.

Custom Protocol Interpreters

Appendix C of the Token Ring manual has been entirely rewritten, and is included as pages 179-200 of TR-Add. In general, it applies equally to the Ethernet and the Token Ring versions.

The only significant addition is that first-level Protocol Interpreters may now be registered for a set of Ethertypes instead of LLC SAPs. To do so, call `register_pi` with `pi_type = PITYPE_ETHTYPE` instead of `PITYPE_SAP`, and provide an array of Ethertypes instead of an array of SAPs.

All other formatting routines and conventions remain the same. Of course, globals such as `llc_header` and `llc_type` have no meaning for Ethertype frames.

Network General Corporation
1296B Lawrence Station Road
Sunnyvale, California 94089

Part No. 20010-1